

METHOD FOR MANIPULATING DATA IN A GROUP OF PROCESSING ELEMENTS TO TRANSPOSE THE DATA

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application is related to the following applications: Method for Manipulating Data in a Group of Processing Elements to Perform a Reflection of the Data (docket no. DB001071-000); Method for Manipulating Data in a Group of Processing Elements (docket no. DB001072-000); Method for Manipulating the Data in a Group of Processing Elements to Transpose the Data Using a Memory Stack (docket no. DB001069-000); Method of Obtaining Interleave Interval for Two Data Values (docket no. DB001065-000); and Method of Rotating Data in a Plurality of Processing Elements (docket no. DB001063-000), all filed concurrently herewith.

BACKGROUND OF INVENTION

[0002] The present invention relates generally to parallel processing and, more specifically, to parallel processing in an active memory device or single instruction, multiple data (SIMD) computer.

[0003] A single, synchronous dynamic random access memory (SDRAM) chip has an internal data bandwidth of greater than 200 G bits/s and a very wide data bus (thousands of bits). That vast data bandwidth provides an opportunity for high performance. Active memories represent one effort to use that vast data bandwidth to improve performance.

[0004] An active memory is a memory device which has a built in processing resource. One of the principal advantages of active memory is that data is processed close to where it is stored. Usually the processing resource is a highly parallel computer system which has processing power to exploit the very high data bandwidths available inside a memory system. An example of an active memory system is illustrated in Figure 1.

[0005] In FIG. 1, a main memory 10 appears as a traditional memory to a CPU 12 except that the main memory 10, by virtue of memory processors 14, can be instructed to perform tasks on its data without the data being transferred to the CPU 12 or to any other part of the system over a

system bus 16. The memory processors 14 are a processing resource distributed throughout the main memory 10. The processing resource is most often partitioned into many similar processing elements (PEs). The PEs are usually simple and operate in parallel. In such a system, the work of the CPU 12 is reduced to various operating system tasks such as scheduling. A substantial portion of the data processing is performed within the main memory 10 by virtue of the memory processors 14.

[0006] Active memory systems have a long history. The earliest systems were built in the 1960's. However, until the advent of integrated logic and current DRAM technologies, active memory computers were always expensive, special machines, excluded from mass market applications. For active memory to be effective, the organization of data in the PE array is an important consideration. Hence, the provision of an efficient mechanism for moving data from one PE to another is an important consideration in the design of the PE array.

[0007] In the past, several different methods of connecting PEs have been used in a variety of geometric arrangements including hypercubes, butterfly networks, one-dimensional strings/rings and two-dimensional meshes. In a two-dimensional mesh or arrays, the PEs are arranged in rows and columns, with each PE being connected to its four neighboring PEs in the rows above and below and columns to either side which are sometimes referred to as north, south, east and west connections.

[0008] Disclosed in G.B. Patent Application Serial No. GB02215 630, entitled Control of Processing Elements in Parallel Processors, filed September 17, 2002 is an arrangement in which a column select line and a row select line can be used to identify processing elements which are active, e.g. capable of transmitting or receiving data. The ability to use a row select signal and a column select signal to identify active PEs provides a substantial advantage over the art in that it enables data to be moved through the array of PEs in a nonuniform manner. However, the need still exists for enabling PEs within the array to work independently of its neighboring PEs even though each PE within the array has received the same instruction.

SUMMARY OF THE INVENTION

[0009] The present invention is directed to a method for transposing data in a plurality of processing elements. The method is comprised of shifting the data along diagonals of the plurality of processing elements until the processing elements in the diagonal have received the

data held by every other processing element in that diagonal. Shifting along diagonals can be accomplished by executing pairs of horizontal and vertical shifts in the x – y directions or pairs of shifts in perpendicular directions, e.g. x-z. Each processing element stores data as its final output data as a function of the processing element's position.

[0010] In one embodiment, an initial count is either loaded into each processing element or calculated locally based on the processing element's location. The initial count may be given by one of the following expressions:

$$\begin{aligned} &(x + y + 1) \text{ MOD (array size)} \\ &(C + R + 1) \text{ MOD (array size)} \\ &(C + y + 1) \text{ MOD (array size) or} \\ &(x + R + 1) \text{ MOD (array size).} \end{aligned}$$

Where X,Y,C,R are defined as shown in Figs 15B and 15C.

Thereafter, the initial count is modified by a programmable amount at programmable intervals, for example, counting down after every pair vertical/horizontal shifts is completed, to produce a current count. When the current count equals a non-zero value, the current data is selected as the final output data. The present invention contemplates hardware, e.g., memory containing an ordered set of instructions, for carrying out the disclosed methods. The present invention provides an efficient method for obtaining the reflection of data. Those advantages and benefits, and others, will be apparent from the description of the invention appearing below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] For the present invention to be easily understood and readily practiced, the present invention will be described in conjunction with an exemplary embodiment, for purposes of illustration and not limitation, in conjunction with the following figures wherein:

[0012] FIG. 1 is a block diagram illustrating the concept of active memory;

[0013] FIG. 2 is a high level block diagram of one example of an active memory on which the methods of the present invention may be practiced;

[0014] FIG. 3 is a high level block diagram of one example of a PE;

[0015] FIG. 4 is a diagram illustrating one type of logic circuit that may be used to interconnect the PE illustrated in FIG. 3 to other PEs;

[0016] FIG. 5 illustrates one method of interconnecting PEs to form an array of PEs;

[0017] FIGs. 6A and 6B illustrate one example of an edge shift;
[0018] FIGs. 7A and 7B illustrate one example of a planar shift;
[0019] FIGs. 8A and 8B illustrate one example of a wrap shift;
[0020] FIGs. 9A and 9B illustrate one example of a vector shift;
[0021] FIGs. 10A and 10B illustrate another example of a vector shift;
[0022] FIGs. 11A and 11B illustrate one example of a data broadcast from the edge registers in which a row and column select function enabled;
[0023] FIGs. 12A and 12B illustrate one example of a broadcast in which only one column is selected;
[0024] FIGs. 13A and 13B illustrate one example of selected edge registers being loaded with the AND of selected columns;
[0025] FIGs. 14A and 14B illustrate another example of a data broadcast;
[0026] FIG. 15A illustrates an initial matrix of data while FIGs. 15B and 15C illustrate transpositions of the data in the matrix of FIG. 15A;
[0027] FIG. 16A illustrates another matrix of data and FIG. 16B illustrates the transpose of that data;
[0028] FIG. 17A illustrates the data appearing in the first row of the matrix of FIG. 16A while FIGs. 17B – 17H illustrate the data for seven consecutive pairs of horizontal/vertical shifts, respectively;
[0029] FIG. 18 illustrates the initial value for the counter within each processing element to effect a transpose of the initial matrix of data; and
[0030] FIGs. 19A – 19L illustrate the steps carried out on the first row of data to effect the transpose.

DESCRIPTION OF THE INVENTION

[0031] Illustrated in FIG. 2 is a high level block diagram of one example of an active memory device 18 on which the methods of the present invention may be practiced. The reader should understand that the methods of the present invention are generally applicable to any group of processing elements having the necessary physical connections between PEs to enable the manipulation of data as required by the methods. The hardware illustrated in FIG. 2 is disclosed for purposes of illustration and not limitation. Furthermore, those of ordinary skill in the art will

recognize that the block diagram of FIG. 2 is an overview of an active memory device with a number of components known in the art being omitted for purposes of clarity.

[0032] The active memory device 18 of FIG. 2 is intended to be deployed in a computer system as a slave device, where a host processor (e.g. CPU 12 in FIG. 1) sends commands to the active memory device 18 to initiate processing within the active memory device 18. A complete processing operation, i.e., data movement and processing, in the active memory device 18 will usually consist of a sequence of many commands from the host to the active memory device 18.

[0033] The active memory device 18 may have two interfaces, a bus interface 20 and a host memory interface 22, for interfacing with the host or other external logic for the purposes of data input, data output and for control. The host memory interface 22 (data input/output ports) of the active memory device 18 is similar in its operation to the interface of a synchronous DRAM. To access data within a DRAM array 24, the host must first activate a page of data. Each page may contain 1024 bytes of data and there may be 16384 pages in all. Once a page has been activated, it can be written and read through the, for example, 32 bit data input/output ports. The data in the DRAM array 24 is updated when the page is deactivated.

[0034] In the active memory device 18 the input and output ports are separate, or they may be combined into a single bi-directional input/output port. A control output may be provided to control a bi-directional buffer servicing the single bi-directional input/output port.

[0035] The host memory interface 22 may operate at twice the frequency of the master input clock. A copy of the 2x clock may be driven off-chip as a timing reference. Unlike a traditional DRAM, the access time for the host memory interface 22 port takes a variable number of cycles to complete an internal operation, such as an activate or deactivate. A ready signal (rdy) is provided to allow the host to detect when the command has been completed.

[0036] The control or command port (cmd) may be a straightforward 32 bit synchronous write/read interface. Writes place both data and the corresponding address into a FIFO 26 of a task dispatch unit 28, which holds the commands until they are executed in the order they were issued. This arrangement allows a burst of commands to be written to the active memory device 18 suiting the burst operation of many peripheral buses. Reads may operate directly.

[0037] The command port is also synchronous, running at the same frequency as the master input clock. Similarly to the host memory interface 22 port, the clock may be driven out as a timing reference.

[0038] In addition to the two address-mapped ports, the active memory device 18 has an interrupt output (intr) which is used to alert its host to various different conditions.

[0039] Internal control in the active memory device 18 is handled by three processors. The task dispatch unit 28 (mentioned above) receives commands from the command port, interprets them and passes them on to the other two processors, a DRAM control unit 30 and an array sequence controller 32. The task dispatch unit 28 also maintains the addresses for operand placement in an array processor register file (RF) and enables access to on-chip resources such as a program memory 34.

[0040] The DRAM control unit 30 controls the DRAM array 24. The DRAM control unit 30 arbitrates between requests for DRAM array 24 access from the host through host memory access registers (H) and through the host memory interface 22. The DRAM control unit 30 also schedules DRAM array 24 refreshes.

[0041] The array sequence controller 32 controls an array or two dimensional mesh of PEs 36. The sequence controller 32 also executes a program from the program memory 34 and broadcasts control signals into the array of PEs 36. The DRAM control unit 30 and array sequence controller 32 may have a synchronization mechanism, whereby they can link the execution of tasks in either processor.

[0042] The active memory device 18 may contain, according to one embodiment, sixteen 64k x128 eDRAM cores. Each eDRAM core is closely connected to an array of sixteen PEs, making 256 (16 x 16) PEs in all.

[0043] FIG. 3 is a high level block diagram of one example of a PE 37. The PE 37 is comprised of a set of Q registers and a shift network 38 which interact with a set of M registers and another shift register 40. One of the sets of registers and shift network 38, 40 receives inputs from various registers, such as register R0, R1, R2 and 0. The output of the registers and shift networks 38, 40 is input to an arithmetic logic unit (ALU) 42. The ALU 42 is capable of performing various arithmetic functions on its input such as addition, subtraction, etc. as is known. The ALU 42 is in communication with condition logic 44 and a result pipe 46.

[0044] The result pipe 46 is a series of interconnected registers R0, R1, R2 and a neighborhood connection register X, which may be used to output a final value. The result pipe 46 also receives through a multiplexer 47 data in the form of an output signal X from its four neighbors, one to the north (XN), one to the east (XE), one to the south (XS) and one to the west (XW). If

the PE 37 happens to be located on an edge of an array, then it may be receiving data from an edge register or a PE in the same row or column, but on an opposite edge, as will be described in greater detail below.

[0045] The result pipe 46 is in communication with a register file (RF) 48 which in turn is in communication with an interface 50. The interface 50 may include a DRAM interface 52 as well as access to the host memory access registers (H).

[0046] The reader should recognize that the PE 37 illustrated in FIG. 3 is exemplary only and is not intended to limit the present invention. For example, the number and location of registers and shift networks may vary, the complexity of the ALU 42 and condition logic 44 may vary, the number of registers and interconnection of registers in the result pipe 46, the size and number of register files, and connection to neighboring PEs as well as other logic may be varied while remaining within the scope of the present invention. The particular architecture illustrated in FIG. 3 was selected to provide a rich register set to enable fairly complex multi-byte operations to be kept within the PE as much as possible.

[0047] For example, the Q registers and shift network 38 allow for data shifting within the 32 bits of the Q register to the left (most significant direction) one, two, four or eight places and eight places to the right as well as for merging data back into a floating point format. The M registers and shift network 40 allow for data shifting within the 32 bits of the M register to the right (least significant direction) one, two, four or eight places and for demerging data from floating point into a signed magnitude plus exponent format. The result from the ALU 42 can be loaded into register R0 while the input from the register file 48 can be loaded into either register R1 or register R2. The neighborhood connection register X can be used as a flexible member of the result pipeline allowing a pipeline length of up to four to be programmed within the PE 37. The X register can be loaded from the R0, R1 or R2 registers, or from the neighborhood interconnection input (the X register of a neighboring PE). The output of the X register can be fed back into the result pipeline at R1 or R2. The register file 48 may be implemented as a 128 entry by 8 bit register file implemented as a synchronous static RAM.

[0048] The DRAM interface 52 may contain two registers, a RAM IN register and a RAM OUT register. Input from the DRAM 24 of FIG. 2 may be held in the RAM IN register while output to the DRAM 24 is held in the RAM OUT register. The RAM IN and RAM OUT registers may reside in the clock domain of the DRAM 24 which typically uses a slower or divided clock

derived from the same source as the clock used for the PE array 36. The RAM IN and RAM OUT registers may be controlled directly from the DRAM control unit 30 and are not visible to the programmer. Data can be transferred into and out of the register file 48 using stolen cycles. Data can also be transferred to/from the host memory access registers (H) without stealing cycles from processing in the PE 37.

[0049] Eight host memory access registers (H) may be provided which allows for a short burst of four or eight bytes to be transferred into or out of the DRAM 24 for host access. Those registers may be multiplexed and be visible from the host memory interface 22 (see FIG. 1) as a page of data. More details about the PEs may be found in G.B. Patent Application No. _____ entitled Host Memory Interface for a Parallel Processor and filed September 17, 2002, which is hereby incorporated by reference.

[0050] FIG. 4 is a diagram illustrating one type of logic circuit that may be used to interconnect PEs of the type illustrated in FIG. 3. The reader will understand that many types of logic circuits may be used to interconnect PEs depending upon the functions to be performed. Using the logic circuit of FIG. 4 to interconnect PEs may result in an array of PEs 36 of the type illustrated in FIG. 5.

[0051] Turning now to FIG. 5, the X register within the result pipe 46 of each PE is driven out as, for example, an eight bit wide X output. Eight bits has been chosen in connection with this architecture as the data width for the PE-PE interconnect to keep a balance between the data movement performance of the array and the improved computational performance. Other sizes of interconnects may be used. The X output is connected to the neighboring inputs of each PE's closest neighbors in the north and west directions. To the south and east, the X output is combined with the input from the opposite direction and driven out to the neighboring PE.

[0052] At the edges of the array 36, the out-of-array connection is selected though a multiplexer to be either the output from the opposite side of the array or an edge/row register 54 or an edge/col. register 56. The edge registers 54, 56 can be loaded from the array output or from the controller data bus. A data shift in the array can be performed by loading the X register from one of the four neighboring directions. The contents of the X register can be conditionally loaded on the AND gate of the row select and column select signals which intersect at each PE. When the contents of the X register is conditionally loaded, the edge registers 54, 56 are also loaded conditionally depending on the value of the select line which runs in the same direction. Hence,

an edge/row register 54 is loaded if the column select for that column is set to 1 and an edge/col register 56 is set if the row select is set to 1. The reader desiring more information about the hardware configuration illustrated in FIG. 5 is directed to G.B. Patent Application GB02215 630, entitled Control of Processing Elements in Parallel Processors filed September 17, 2002, which is hereby incorporated by reference.

[0053] With the hardware previously described, a number of shifting operations may be performed as illustrated in FIGs. 6A, 6B through 10A, 10B. In FIG. 6A and 6B, an edge shift is illustrated. In the edge shift, the edge/col registers 56 are active as the data is shifted left to right (west to east) as shown in FIGs. 6A, 6B. The reader will recognize that an edge shift may be performed in the other direction, right to left (east to west). Alternatively, edge shifts may be performed by using the edge/row registers 54 in a north to south or south to north direction.

[0054] Illustrated in FIGs. 7A, 7B is a planer shift. In the planer shift there is no wrap around from the edge of the array. The reader will recognize that in addition to the planer shift illustrated in FIGs. 7A, 7B, planer shifts from east to west, north to south, and south to north may also be performed.

[0055] Illustrated in FIGs. 8A, 8B is a wrap shift. In the wrap shift, the edge/col registers 56 do not participate. Additionally, wrap shifts from east to west, north to south and south to north may be performed.

[0056] Illustrated in FIGs. 9A, 9B is a vector shift. Again, the edge/col registers 56 do not participate. Furthermore, the output of the PE in the bottom right corner of the array wraps to the input of the PE in the upper left corner of the array. In FIGs. 10A and 10B, a vector shift in the direction opposite to the direction of FIGs. 9A, 9B is illustrated. The reader will recognize that vector shifts from north to south and south to north may also be performed.

[0057] Returning to FIG. 5, the PE-PE interconnect may also provide a broadcast and broadcast network. Connections or buses 58 extend north to south from a column select register 59 and connections or buses 60 extend west to east from a row select register 61. Also provided is row broadcast/broadcast AND chain 62 and a column broadcast/broadcast AND chain. When used for data broadcast or broadcast, these connections (column buses 58 and row buses 60) act as if driven by open drain drivers; the value on any bit is the wire-AND of all the drivers outputs. Three control signals (broadcast, broadcast and intercast) determine the direction of the buses as follows:

- If broadcast is set to 1, any PE for which the corresponding bits of the row select register 61 and column select register 59 are both set will drive both the row buses 60 and the column buses 58. Note that if no PEs in a row or column drive the bus, the edge register at the end of that row or column will be loaded with 0 x FF.
- If broadcast is set to 1, the row bus 60 is driven from the row select register 61 and the column bus 58 is driven from the column select register 59 and any PE for which the corresponding bits of the row select register 61 and column select register 59 are both set will be loaded from one of the row or column inputs, according to which is selected.
- If intercast is set to 1, any PE in which its A register is 1 will drive its output onto its row bus 60 and column bus 58 and any PE for which the corresponding bits of the row select register 61 and column select register 59 are both set will be loaded from one of the row buses 60 or column buses 58, according to which is selected.

[0058] With the aforementioned connections, a number of operations are possible, some of which are illustrated in FIGs. 11A, 11B through 14A, 14B.

[0059] In FIGs. 11A, 11B, data is broadcast from edge/col registers 56 with the column select register 59 and row select register 61 set as illustrated in FIG. 11A. As a result, data is latched in the PEs as shown in FIG. 11B in which four PEs are active, and the remainder of the PEs in the array are inactive.

[0060] In FIGs. 12A, 12B, a broadcast instruction is illustrated in which one column is selected by setting the value for that column's bus 58 to 1. In this broadcast-column operation, only those edge /col. registers 56 for which the row select register 61 bits are set, will be loaded. Similarly, in a broadcast-row operation (not shown), only those row/edge registers 54 for which the corresponding column select register 59 bits are set, will be loaded.

[0061] FIGs. 13A, 13B illustrate a broad catch instruction. In the illustrated example, the column select register 59 and row select register 61 are used to select the PEs whose values will be AND'ed together and loaded in the corresponding edge/col registers 56. In FIGs. 13A, 13B, the column edge registers 56 are loaded with the AND of selected columns, except where the row select is 0.

[0062] In FIGs. 14A, 14B, an example of an intercast operation is illustrated. In an intercast operation, PEs which drive onto the row buses 60 and column buses 58 are determined by each

PE's A register value. The PEs which are loaded are determined by the row and column selects, just like for a broadcast. In FIG. 14A, data is broadcast from the X registers of those PEs where A equals 1 while in FIG. 14B, the column select register 59 and row select register 61 together activate those PEs into which data will be written.

[0063] Using the aforementioned instructions or operations, a group of instructions may be combined into an instruction set for manipulating data within the array 36 of PEs. The instruction set may include a single instruction or operation or a combination of instructions. Each individual instruction is carried out though a series of shifts.

[0064] In operation, an input matrix of data is placed on the shift network, and moved around by using a combination of north, south, east and west shifts. In addition, the column select register 59 and row select register 61 may be used to determine which of the PEs is active. The exact combination of active PEs, instructions, and direction in which the instruction (shift) is performed will depend upon the particular array manipulation required. As the instructions are executed and the shifting proceeds, each PE will be presented with different array values. For example, if a wrap shift is performed a number of times equal to the number of PEs in a row, each PE in the row will see every value held by all of the other PEs in the row.

[0065] A PE can conditionally select any of the values it sees as its final output value by conditionally loading that value, which is representative of an output result matrix. However, only one value, the desired result, is loaded.

[0066] All X values are passed through the PE; the required output value is conditionally loaded once it has arrived in the PE. The conditional loading can be done in various ways. e.g. by using any PE registers except X, R1, or R2. An example is shown below.

Clock	PE C + 0		PE C + 1		PE C + 2		PE C + 3
Cycle							
T + 0	X <= xe(east)	←	X <= xe	←	X <= xe	←	X <= xe
	↓		↓		↓		↓
T + 1	R1 <= X		R1 <= X		R1 <= X		R1 <= X
	↓		↓		↓		↓
T + 2	<cond>?R0 <= R1		<cond>?R0 <= R1		<cond>?R0 <= R1		<cond>?R0 <= R1

- At time T+0: The X register reads data from the X register on the PE to the East. This shifts data to the left (or West).
- At time T+1: The R1 register unconditionally reads the data off the shift network (X register)
- At time T+2: The R0 register conditionally loads the data from R1. (i.e. if <cond>=1).

[0067] The timing of the loading is achieved by maintaining a current count in a local counter, which is typically implemented in software. In one embodiment, the local counter is set to an initial value. The local counter can be set in a variety of ways, including loading the counter with the initial value or calculating the initial value locally based on the processing element's location in the matrix (or array) and the function being performed on the data. Thereafter, at certain points in the shifting process, the counter is decremented. For example, the counter may be decremented once for each shift that occurs, or may be decremented once per n clock cycles where n clock cycles equals one shift. As stated, the initial value of the counter depends on its position in the matrix or array and is given by the general function $f(\text{Row_Index}, \text{Col_Index})$, where the exact form of $f()$ will depend on the particular array manipulation required. When the counter reaches a non-positive value (i.e., zero or negative) the PE selects the data to be loaded into the output matrix.

[0068] Other ways of achieving the same result include resetting the counter to zero and loading each PE with a target value. Thereafter, the counter is incremented producing a current count. When the current count equals the target value, the data value is selected as the final output value to be loaded into the output matrix. Generally, a counter is set to a first known value. Then, at certain programmable points in the algorithm, the value of the counter may be altered, up or down, by a programmable amount. Storing, occurs when a current count in the counter hits a pre-defined target value.

[0069] FIG. 15A illustrates a matrix of data while FIGs. 15B and 15C illustrate two possible transposes of the matrix of data illustrated in FIG. 15A. In each case, the leading diagonals (shown in bold) remain unchanged. A matrix of data may be represented by PEs arranged in rows and columns extending in the x and y directions, respectively, or PEs arranged in higher dimensional structures, but with perpendicular 'rows' and 'columns', e.g. x and z directions.

[0070] The two examples of FIGs. 15B and 15C can be viewed as using different representations for row order. The present invention need not comply to any conventions for it to work; the present invention need only be internally consistent. The manner in which the rows and columns are numbered is the manner which is quickest to implement. The numbering conventions for possible versions of the algorithm are as follows. For the ‘top’ mode (FIG. 15B), the data can be transposed by moving data across the array in one of two directions:

- right and upwards (or vice versa) using $(X+Y+1) \text{ MOD size}$
- left and downwards (or vice versa) using $(C+R+1) \text{ MOD size}$

For the ‘bottom’ mode (FIG. 15C), the data can be transposed by moving data across the array in one of two directions:

- left and upwards (or vice versa) using $(C+Y+1) \text{ MOD size}$
- right and downwards (or vice versa) using $(X+R+1) \text{ MOD size}$

[0071] Turning to FIGs. 16A and 16B, FIG. 16B represents one possible transpose of the data of FIG. 16A. FIG. 17A illustrates the data appearing in the first row of the matrix of FIG. 16A while FIGs. 17B through 17H illustrate the data as it appears as a result of seven consecutive horizontal/vertical pairs of shifts, i.e., move data right (east) 1 PE, and move data up (north) one PE. In FIG. 17A, the data “a1” appearing in position one represents that data in the first row of the transpose of FIG. 16B. So that data is selected as the final output. In FIG. 17B, after the first pair of east/north shifts, the data “a2” appearing in position two represents the data in the first row of the transpose of FIG. 16B. Similarly, in FIGs. 17C through 17H the data “a3” to “a8” appears in positions three through eight, respectively. Using the expression $(X+Y+1) \text{ MOD } (8)$, the local counters in each of the processing elements may be set to an initial count as illustrated in FIG. 18.

[0072] FIGs. 19A through 19L illustrate input data, output data and counter values through the process of shifting data. In FIG. 19A, the input data is initialized and the counters set to their initial values as illustrated. Thereafter, as illustrated in FIG. 19B, the counter values are tested and the counter in position one is nonpositive. That PE’s current data (“a1”) is selected as the final output data. In FIG. 19C the counters are decremented resulting in the values shown in the third row of FIG. 19C. Thereafter, in step 4, the first pair of shifts is performed resulting in the values shown in the first row of FIG. 19D.

[0073] In FIGs. 19E, 19F, and 19G, the process is repeated. Beginning with FIG. 19E, the counter values for PEs which have not yet latched data are tested, and it is determined that the counter for the PE in position two is not positive. Accordingly, the data “a2” is latched. Thereafter, in step 19F, each counter is decremented again resulting in the values shown in the third row in FIG. 19F. In FIG. 19G, the second pair of shifts is performed resulting in the data illustrated in the first row of FIG. 19G.

[0074] At FIG. 19H, the process is repeated beginning with the testing of the counters for those PEs which have not yet latched data. In FIG. 19H, it is determined that the local counter in the PE in position three is nonpositive such that the data “a3” is latched as the final output data. Thereafter, in FIG. 19I, the value of the counters is once again decremented resulting in the current counts as shown in the third row of FIG. 19I.

[0075] The process is continually repeated until step 22, represented in FIG. 19J, which illustrates the seventh pair of shifts. In FIG. 19K, the counters which have not yet latched data are tested and it is determined that the current count in the processing element in position eight is nonpositive. Accordingly, the data “a8” is latched by the processing element in position eight. In step 23, represented in FIG. 19L, the counters are decremented once last time, although that step is not necessary as it can be seen from FIG. 19K, second row, that all of the data needed for a transpose of the first row has been latched. Figures 19A through 19L illustrate one embodiment for maintaining a current count in local counters within each processing element. Those of ordinary skill in the art will recognize that other methods of maintaining a current count, as well as counting up and counting down, testing for values other than nonpositive values, e.g. a target value, can be carried out while remaining within the scope of the present invention.

[0076] By using the method of the present invention, PEs within a group of PEs can be individually controlled as to the output value which the PE selects for output into the final matrix. Thus, although all of the PEs are responding to the same command, e.g., an east to west wrap shift, each of the PEs is capable of selecting different data at different points during the execution of the instruction thereby enabling various types of data manipulations. Furthermore, by determining which PEs are active, additional flexibility is provided so that subsets of data can be manipulated.

[0077] Although the figures illustrate a two-dimensional (2D) array connected as a mesh the present invention is applicable to other configurations. Accordingly, the phrase “plurality of processing elements” is to be broadly construed to include 2D and 3D collections of PEs connected in any known manner. For example, the PE’s could be connected in shapes other than as illustrated in the figures, e.g., a cube. That would have $f(x_Index, y_Index, z_Index)$. An n-dimensional hypercube would have n dimensions and $f(d(0), d(1), d(2) \dots d(n-1))$.

[0078] Additionally, the network need not be connected as a mesh. For example, a simple extension may be implemented by providing two extra connections, one to the PE halfway across the row and the other to the PE halfway down the column. For that example there would be two more shift connections. In addition to the north, east, south, and west shifts, there could also be Half_Row and Half_Col shifts. Both of the above changes could be used at the same time. For example, a four dimensional hyper-cube with half-way connections would have twelve shift options. However, there are only certain configurations of PE’s in which a transpose operation makes geometric sense. Consider a 3D cuboid of $8 \times 8 \times 4$ PE’s. The 3D transpose would be performed using a “plane of diagonals”. There are 4 slices of 8×8 PE’s, within each slice, the diagonal is used as described above to perform the transpose. Within the cuboid, these diagonals on each of the 4 slices sit on top of each other to form a 4×8 plane of PE’s. In general an N-dimensional mesh of PE’s that can be ‘sliced’ to produce a set of $M \times M$ squares is amenable to a N-dimensional transpose. Furthermore, certain such meshes may be sliced in different directions to produce different sets of squares, e.g., an $8 \times 8 \times 4 \times 4$ mesh could be sliced to produce 8×8 or 4×4 squares.

[0079] While the present invention has been described in connection with a preferred embodiment thereof, those of ordinary skill in the art will recognize that many modifications and variations such as those previously discussed are possible. The present invention is not to be limited by the foregoing description but only by the following claims.